
A NOTE ON THE PARTICLE FLOW FILTER/GENERATOR

Liu Yang

Division of Applied Mathematics
Brown University
Providence, RI 02912, USA
liu.yang@brown.edu

ABSTRACT

We shall analyze particle flow filter in this note.

1 INTRODUCTION

Particle flow filters are a family algorithms developed to sample from distributions of interest, for example, posterior. It's basically trying to do the same thing as MCMC, which has the following disadvantages:

- MCMC methods are almost always computationally expensive in high-dimensional state spaces.
- MCMC methods are highly sequential, which makes it hard for parallel computing.
- It's not convenient to draw independent samples from MCMC.

Particle flow filters tackle this problem by looking for an ODE or SDE governing the motion of particles, so that the distribution of the particles start from the source distribution, e.g. prior, and end up with the target distribution, e.g. posterior.

Once we find the ODE/SDE, we can have independent samples from the target distribution by:

- Step 1: independently sample from the source distribution as particles.
- Step 2: solve ODE/SDE with the starting point of each particle as the initial condition.

Note that there is no interactions between particles, therefore this particle flow approach is highly parallelizable.

Actually there are too many ways to transport from the source distribution to the target distribution. To get the governing ODE/SDE:

- We need to assume a trajectory of distributions, which start from the source distribution and end at the target distribution.
- And the ODE/SDE also yields a trajectory of distributions, given by Fokker-Planck Equation.
- The two trajectories should overlap, so we have an equation for the drift term and diffusion term in ODE/SDE.

Note that the trajectory is not unique. It's actually a good question to ask, what is the optimal trajectory, in what sense.

2 DETERMINISTIC PARTICLE FLOW FILTER

2.1 GENERAL FORMULATION

Consider log-homotopy based particle flow:

$$\log p(\mathbf{x}, t) = \log g(\mathbf{x}) + t \log h(\mathbf{x}) - \log K(t), \quad (1)$$

where $g(\mathbf{x})$ and $h(\mathbf{x})$ are both probability density functions, and $K(t)$ is the normalization term so that $p(\mathbf{x}, t)$ is a probability density function for all $t \in [0, 1]$. Note that $\log K(0) = 0$. In the context of Bayesian learning, g is the density of prior, h is the density of likelihood, so that $p(\cdot, 1)$ is the density of posterior.

For deterministic particle flow

$$\frac{d\mathbf{u}}{dt} = \mathbf{v}(\mathbf{u}, t), \quad (2)$$

we have Fokker-Planck equation with zero diffusion:

$$\frac{\partial p(\mathbf{x}, t)}{\partial t} = -\nabla \cdot (\mathbf{v}p). \quad (3)$$

Then we have

$$\begin{aligned} \nabla \cdot (\mathbf{v}p) &= -p \frac{\partial \log p}{\partial t} \\ &= -p(\log h - \frac{\partial}{\partial t} \log K(t)) \end{aligned} \quad (4)$$

i.e.

$$\nabla \cdot \mathbf{v} + \mathbf{v} \cdot \nabla \log p = -\log h + \frac{\partial}{\partial t} \log K(t) \quad (5)$$

Further, if we assume the particle flow is divergence free, i.e.,

$$\nabla \cdot \mathbf{v} = 0, \quad (6)$$

then we have

$$\mathbf{v} \cdot \nabla \log p = -\log h + \frac{\partial}{\partial t} \log K(t) \quad (7)$$

If we construct a divergence free velocity field \mathbf{v} satisfying Equation 7, then starting from density g , we will get to density gh with trajectory $p(\mathbf{x}, t)$.

2.2 FRED'S FORMULATION

In Fred's formulation (Daum, 2016), \mathbf{v} is given by

$$\mathbf{v} = -\frac{\nabla \log p}{\|\nabla \log p\|^2} \log h \quad (8)$$

Which leads to

$$\frac{\partial}{\partial t} \log K(t) = 0 \quad (9)$$

Note that $\log K(0) = 0$, we have

$$\log K(t) = 0, \quad \forall t \in [0, 1] \quad (10)$$

i.e.

$$p(\mathbf{x}, t) = g(\mathbf{x})h^t(\mathbf{x}). \quad (11)$$

And $g(\mathbf{x})h^t(\mathbf{x})$ has to be a probability density function. This is too strong from my point of view. Also I cannot see why $\nabla \cdot \mathbf{v} = 0$ in the above formulation.

2.3 NECESSARY CONDITION FOR THE EXISTENCE OF DIVERGENCE FREE FLOW

Actually, if the velocity field is divergence free, p and f are limited to a very small family of distributions.

Note that the Fokker-Planck Equation has an equivalent form

$$\frac{d}{dt} \log p(\mathbf{u}(\mathbf{x}_0, t), t) = -\nabla \cdot \mathbf{v}(\mathbf{u}, t), \quad (12)$$

where $\mathbf{u}(\mathbf{x}_0, \cdot)$ is the trajectory of the particle start from \mathbf{x}_0 . Since \mathbf{v} is divergence free, we then have

$$\frac{d}{dt} \log p(\mathbf{u}(\mathbf{x}_0, t), t) = 0 \quad (13)$$

That is to say, the density p is constant along the particle trajectory.

3 STOCHASTIC PARTICLE FLOW FILTER

3.1 GENERAL FORMULATION

Consider stochastic particle flow

$$d\mathbf{u} = \mathbf{v}(\mathbf{u}, t)dt + \boldsymbol{\sigma}(\mathbf{u}, t)d\mathbf{W}(t), \quad (14)$$

where \mathbf{v} is the drift term, $\boldsymbol{\sigma}$ is the spatial homogeneous diffusion term (for simplicity). We have Fokker-Planck equation:

$$\frac{\partial p(\mathbf{x}, t)}{\partial t} = -\nabla \cdot (\mathbf{v}p) + \nabla \cdot (\mathbf{Q}\nabla p), \quad (15)$$

where $\mathbf{Q} = \frac{1}{2}\boldsymbol{\sigma}\boldsymbol{\sigma}^T$ is the diffusion tensor. Then we have:

$$\begin{cases} \log p(\mathbf{x}, t) = \log g(\mathbf{x}) + t \log h(\mathbf{x}) - \log K(t) & \text{(i)} \\ \frac{\partial p(\mathbf{x}, t)}{\partial t} = -\nabla \cdot (\mathbf{v}p) + \nabla \cdot (\mathbf{Q}\nabla p) & \text{(ii)} \end{cases} \quad (16)$$

$$\text{(i)} \Rightarrow \frac{\partial \log p(\mathbf{x}, t)}{\partial t} = \log h(\mathbf{x}) - \frac{d \log K(t)}{dt}$$

$$\text{(ii)} \Rightarrow \frac{\partial \log p(\mathbf{x}, t)}{\partial t} = -\nabla \cdot \mathbf{v} - \mathbf{v} \cdot \nabla \log p + \text{Tr}[\mathbf{Q}H(\log p)] + \nabla \log p \cdot \mathbf{Q} \cdot \nabla \log p$$

where H is the Hessian matrix. So we have

$$\log h(\mathbf{x}) - \frac{d \log K(t)}{dt} = -\nabla \cdot \mathbf{v} - \mathbf{v} \cdot \nabla \log p + \text{Tr}[\mathbf{Q}H(\log p)] + \nabla \log p \cdot \mathbf{Q} \cdot \nabla \log p \quad (17)$$

3.2 NORMALIZATION TERM

Now let's deal with the normalization term. Note that:

$$K(t) = \int g(\mathbf{x})h(\mathbf{x})^t d\mathbf{x} \quad (18)$$

Therefore

$$\begin{aligned} \frac{d \log K(t)}{dt} &= \frac{1}{K(t)} \frac{dK(t)}{dt} \\ &= \frac{\int g(\mathbf{x})h(\mathbf{x})^t \log h(\mathbf{x}) d\mathbf{x}}{\int g(\mathbf{x})h(\mathbf{x})^t d\mathbf{x}} \\ &= \mathbb{E}_{p(\mathbf{x}, t)} \log h(\mathbf{x}), \end{aligned} \quad (19)$$

where $\mathbb{E}_{p(\mathbf{x}, t)}$ means expectation with density $p(\mathbf{x}, t)$.

Remark: $K(t)$ is the partition function of $p(\mathbf{x}, t)$, so the derivative $d \log K(t)/dt$ should be the $p(\mathbf{x}, t)$ -expectation of coefficient of t in $\log p$, which is $\log(h)$ here.

3.3 GOVERNING EQUATION

Combining the equations we have the governing equation for the dynamics:

$$\log h(\mathbf{x}) - \mathbb{E}_{p(\mathbf{x}, t)} \log h(\mathbf{x}) = -\nabla \cdot \mathbf{v} - \mathbf{v} \cdot \nabla \log p + \text{Tr}[\mathbf{Q}H(\log p)] + \nabla \log p \cdot \mathbf{Q} \cdot \nabla \log p \quad (20)$$

Note the following:

- There is no normalization term in derivatives of $\log p$, i.e. we have exact analytic expression of $\nabla \log p$ and $H(\log p)$.
- $p(\mathbf{x}, t)$ -expectation could be approximated empirically with particles at time t as samples of $p(\mathbf{x}, t)$. A numerical approach: solve \mathbf{v} and \mathbf{Q} until time t_k , then go to t_{k+1} .
- In general, solving this equation is not easy, if g and h are complicated. But maybe we can use techniques from deep learning.

3.4 NUMERICAL EXPERIMENT

Using neural networks to solve v .

- Prior: Gaussian centered at $(1, 1)$, var = 1.0.
- Likelihood: Gaussian centered at $(-1, -1)$, var = 1.0.
- Posterior: Gaussian centered at $(0, 0)$, var= 0.5.




Figure 1: Left: map from $u(0)$ to $u(1)$. Right: samples of prior (blue) and posterior (orange). Mean of posterior samples: $(-0.00346315, -0.00049885)$, variance of posterior samples: $(0.49337277, 0.50254256)$. 10^6 samples.

3.5 NEGLECTING NORMALIZATION TERM

In Drum Fred's formulation, the authors start from Equation 17, and take gradient of x to remove the normalization term $K(t)$, which is equivalent to

$$\begin{aligned} \log h(\mathbf{x}) + C(t) = & -\nabla \cdot \mathbf{v} - \mathbf{v} \cdot \nabla \log p + \text{Tr}[\mathbf{Q}H(\log p)] \\ & + \nabla \log p \cdot \mathbf{Q} \cdot \nabla \log p \end{aligned}$$

I actually got puzzled by the following paradox:

- If we neglect $K(t)$, then $p(\mathbf{x}, t)$ could be unnormalized, i.e. $\int_{\mathbf{x}} p(\mathbf{x}, t) d\mathbf{x} \neq 1$.
- On the other hand, $p(\mathbf{x}, t)$ represents the density of the particles. The particles never vanish, how could $\int_{\mathbf{x}} p(\mathbf{x}, t) d\mathbf{x} \neq 1$?

But later I figured out, it's about the vanishing condition of $\mathbf{v}p$ and $\mathbf{Q}\nabla p$ at infinity. Think about this 1D toy case:

$$\begin{aligned} p(x, 0) &= \mathbf{1}_{x \geq 0} \frac{1}{(x+1)^2} \\ p(x, 1) &= \mathbf{1}_{x \geq 0} \frac{1}{(2x+1)^2} \end{aligned}$$

We can find a deterministic flow for p above ($x \rightarrow 2x$), but $\mathbf{v}p$ (flux) won't vanish at infinity. Actually, the time integral of flux converges to 0.5 as x goes to infinity, which means v goes to infinity.

Theoretically, it doesn't matter if p is normalized or not, since we only care about samples. But I think the explosion of v may lead to some numerical difficulty. So I would prefer not to neglect the normalization term.

4 CONCLUSION

Particle flow filter can be used to sample from (unnormalized) density. The time derivative of the normalization term could be formulated as an expectation, and be numerically tackled via sampling.

Solving the governing equation of particle flow is difficult in general, but it's possible to apply deep learning techniques. If we neglect the normalization term, theoretically it doesn't matter since we only care about samples. But the flux won't vanish at infinity, which might be dangerous from a numerical point of view.

REFERENCES

Fred Daum. seven dubious methods to compute optimal q for bayesian stochastic particle flow. In *2016 19th International Conference on Information Fusion (FUSION)*, pp. 2237–2244. IEEE, 2016.

APPENDIX

Code for the numerical experiment.

```
1 import logging, os
2 logging.disable(logging.WARNING)
3 os.environ['TF_CPP_MIN_LOG_LEVEL'] = '3'
4
5
6 import argparse
7 import numpy as np
8 import tensorflow as tf
9 import random
10 import matplotlib.pyplot as plt
11
12 import sys
13
14
15
16 def feed_NN(X, W, b, act = tf.nn.tanh):
17     A = X
18     L = len(W)
19     for i in range(L-1):
20         A = act(tf.matmul(A, W[i]) + b[i])
21     return tf.matmul(A, W[-1]) + b[-1]
22
23
24 def logPden_fun(x, dim = 2):
25     sigma = 1.0
26     logPden = - 0.5 * tf.reduce_sum((x - 1)**2/(sigma**2), axis = 1) \
27         - 0.5 * dim * np.log(2*np.pi) - np.log(sigma) * dim
28     return logPden
29
30 def logQden_fun(x, Qden, eps, dim = 2):
31     sigma = 1.0
32     logQden = - 0.5 * tf.reduce_sum((x + 1)**2/(sigma**2), axis = 1) \
33         - 0.5 * dim * np.log(2*np.pi) - np.log(sigma) * dim
34     return logQden
35
36
37 class PdataGenerator():
38     def __init__(self, allsize = 40000):
39         self.allsize = allsize
40         self.alldata = (np.random.normal(0,1,[allsize,2]) * np.array
41             ([1,1])).astype(np.float32) + 1.0
42
43     def nextbatch(self, bs):
44         indexes = np.random.choice(self.allsize, bs, replace = False)
45         return self.alldata[indexes,:]
46
47     def nextnewbatch(self, bs):
```

```

47     return (np.random.normal(0,1,[self.allsize,2]) * np.array([1,1])).
48     astype(np.float32) + 1.0
49
50
51 def NN(x_in, dim_out, depth, width, act):
52     A = x_in
53     for _ in range(depth):
54         A = tf.layers.dense(A, units= width, activation = act)
55     A = tf.layers.dense(A, units= dim_out, activation = None)
56     return A
57
58
59 def main(args):
60
61     os.environ['CUDA_DEVICE_ORDER']='PCI_BUS_ID' # see issue #152
62     os.environ['CUDA_VISIBLE_DEVICES']= args.GPU
63
64     tf.set_random_seed(args.random_seed)
65     np.random.seed(args.random_seed)
66
67     P = PdataGenerator()
68
69     steps = args.steps; dt = 1.0 / steps
70     Pdata = tf.placeholder(tf.float32, [None,2])
71
72
73
74     if args.act == 'tanh':
75         act = tf.nn.tanh
76     elif args.act == 'lrelu':
77         act = tf.nn.leaky_relu
78     elif args.act == 'relu':
79         act = tf.nn.relu
80     elif args.act == 'softplus':
81         act = tf.nn.softplus
82     else:
83         raise NotImplementedError
84
85     v = [None for i in range(args.steps)]
86     loss = [None for i in range(args.steps)]
87     opt = [None for i in range(args.steps)]
88     us = [None for i in range(args.steps+1)]
89     us[0] = Pdata
90
91     if args.problem == 'Generative':
92         for ti in range(0,args.steps):
93             logPden = logPden_fun(us[ti])
94             logQden = logQden_fun(us[ti], args.Qden, args.eps)
95             gradlogp = (1-ti*dt) * tf.gradients(logPden, us[ti])[0] \
96                 + (ti*dt) * tf.gradients(logQden, us[ti])[0]
97             with tf.variable_scope('v' + str(ti), reuse = tf.AUTO_REUSE):
98                 v[ti] = NN(us[ti], args.dim, args.nn_depth, args.nn_width
99                 , act)
100             divv = tf.reduce_sum([tf.gradients(v[ti][:,i], us[ti])[0][:,i]
101             ] for i in range(args.dim)], axis = 0)
102             us[ti+1] = us[ti] + dt * v[ti]
103             LHS = logQden - logPden - tf.reduce_mean(logQden - logPden,
104             axis=0) # (batch_size,)
105             RHS = - divv - tf.reduce_sum(v[ti] * gradlogp, axis=1) # (
106             batch_size,)
107             loss[ti] = tf.reduce_mean((LHS - RHS)**2)
108             var_list = [i for i in tf.trainable_variables() if 'v' + str(
109             ti) in i.name]

```

```

105     opt[ti] = tf.train.AdamOptimizer(learning_rate=1e-4).minimize
      (loss[ti],
106         var_list=var_list)
107     print(ti, end = ' ', flush=True)
108
109     elif args.problem == 'Bayesian':
110         for ti in range(0, args.steps):
111             logPden = logPden_fun(us[ti])
112             logQden = logQden_fun(us[ti], args.Qden, args.eps)
113             gradlogp = tf.gradients(logPden, us[ti])[0] \
114                 + (ti*dt) * tf.gradients(logQden, us[ti])[0]
115             with tf.variable_scope('v' + str(ti), reuse = tf.AUTO_REUSE):
116                 v[ti] = NN(us[ti], args.dim, args.nn_depth, args.nn_width
, act)
117             divv = tf.reduce_sum([tf.gradients(v[ti][:,i], us[ti])[0][:,i
] for i in range(args.dim)], axis = 0)
118             us[ti+1] = us[ti] + dt * v[ti]
119             LHS = logQden - tf.reduce_mean(logQden, axis=0) # (batch_size
,)
120             RHS = - divv - tf.reduce_sum(v[ti] * gradlogp, axis=1) # (
batch_size,)
121             loss[ti] = tf.reduce_mean((LHS - RHS)**2)
122             var_list = [i for i in tf.trainable_variables() if 'v' + str(
ti) in i.name]
123             opt[ti] = tf.train.AdamOptimizer(learning_rate=1e-4).minimize
      (loss[ti],
124         var_list=var_list)
125             print(ti, end = ' ', flush=True)
126
127     else:
128         raise NotImplementedError
129
130
131     config = tf.ConfigProto()
132     config.gpu_options.allow_growth = True
133     sess = tf.Session(config=config)
134     sess.run(tf.global_variables_initializer())
135
136     step = 0
137
138     savedir = 'save' + \
139         '-Qden'+ args.Qden + \
140         '-flow-' + str(args.flow) + '-' + str(args.nn_depth) + 'x' +
str(args.nn_width) + '-' + str(args.act) + \
141         '-seed' + str(args.random_seed)
142
143     if not os.path.exists(savedir):
144         os.mkdir(savedir)
145     saver = tf.train.Saver(max_to_keep=1000)
146
147     if args.restore > 0:
148         saver.restore(sess, savedir+'/' + str(args.restore) + '.ckpt')
149         P_test = P.nextnewbatch(1000000)
150         thisloss = sess.run(loss, feed_dict = {Pdata: P_test})
151         out = sess.run(us[-1], feed_dict={Pdata: P_test})
152         print(step, np.sum(thisloss), np.mean(out, axis=0), np.var(out,
axis=0), flush = True)
153
154         plt.figure(figsize = (9,4))
155         plt.subplot(1,2,1)
156         thisP = P.nextbatch(args.batch_size)
157         out = sess.run(us[-1], feed_dict={Pdata: thisP})
158         for index in range(100):
159             plt.arrow(thisP[index,0], thisP[index,1],
out[index,0] - thisP[index,0],

```

```

161         out[index,1] - thisP[index,1],
162         head_width=0.1,
163         head_length=0.1, color = 'k')
164     plt.xlim((-3,3))
165     plt.ylim((-3,3))
166
167     plt.subplot(1,2,2)
168     plt.scatter(thisP[:,0], thisP[:,1], s = 5, label = 'input')
169     plt.scatter(out[:,0], out[:,1], s = 5, label = 'output')
170     plt.legend()
171     plt.xlim((-3,3))
172     plt.ylim((-3,3))
173     plt.savefig(savedir + '/' + str(args.restore)+'.eps', format = '
eps')
174     return
175
176     for ti in range(0,args.steps):
177         for i in range(args.iterations + 1):
178             if step % 500 ==0:
179                 saver.save(sess, savedir+'/' + str(i) + '.ckpt')
180             if step % 100 ==0:
181                 plt.figure(figsize = (9,4))
182                 plt.subplot(1,2,1)
183                 thisP = P.nextbatch(args.batch_size)
184                 out = sess.run(us[-1], feed_dict={Pdata: thisP})
185                 for index in range(100):
186                     plt.arrow(thisP[index,0], thisP[index,1],
187                             out[index,0] - thisP[index,0],
188                             out[index,1] - thisP[index,1],
189                             head_width=0.1,
190                             head_length=0.1, color = 'k')
191                 plt.xlim((-3,3))
192                 plt.ylim((-3,3))
193
194                 plt.subplot(1,2,2)
195                 plt.scatter(thisP[:,0], thisP[:,1], s = 5, label = 'input
')
196                 plt.scatter(out[:,0], out[:,1], s = 5, label = 'output')
197                 plt.legend()
198                 plt.xlim((-3,3))
199                 plt.ylim((-3,3))
200
201                 plt.savefig(savedir + '/' + str(step)+'.png', dpi=50)
202
203
204                 # _, thisloss = sess.run([opt[ti], loss[ti]], feed_dict = {
Pdata: P.nextbatch(args.batch_size)})
205                 # print(step, thisloss, flush = True)
206
207                 sess.run(opt, feed_dict = {Pdata: P.nextbatch(args.batch_size
)})
208                 thisloss = sess.run(loss, feed_dict = {Pdata: P.nextbatch(
args.batch_size)})
209                 out = sess.run(us[-1], feed_dict={Pdata: P.nextbatch(args.
batch_size)})
210                 print(step, np.sum(thisloss), np.mean(out, axis=0), np.var(
out, axis=0), flush = True)
211
212                 step += 1
213
214
215
216 if __name__ == '__main__':
217
218     parser = argparse.ArgumentParser(description='PDE Flow Generator')

```



```

219 parser.add_argument('--dim', type = int, default=2,
220                     help='space dimension')
221 parser.add_argument('--GPU', type = str, default='0',
222                     help='GPU index')
223 parser.add_argument('-fl', '--flow', choices=['PDE'], default= 'PDE',
224                     help='model')
225 parser.add_argument('-rs', '--random_seed', type = int, default= 0,
226                     help='random seed for numpy and tensorflow')
227 parser.add_argument('-bs', '--batch_size', type = int, default= 1000,
228                     help='batch size of training')
229 parser.add_argument('--eps', type = float, default= 1e-4,
230                     help='epsilon value used in log function')
231 parser.add_argument('-nd', '--nn_depth', type = int, default= 2,
232                     help='number of hidden layers')
233 parser.add_argument('-nw', '--nn_width', type = int, default= 64,
234                     help='width of hidden layers')
235 parser.add_argument('--steps', type = int, default= 50,
236                     help='steps of ODE flow')
237 parser.add_argument('-it', '--iterations', type = int, default= 1000,
238                     help='steps of ODE flow')
239 parser.add_argument('--Qden', choices=['Eight1', 'Eight2', 'Nine'],
240                     default = 'Nine',
241                     help='density of target distribution Q')
241 parser.add_argument('--act', choices=['tanh', 'lrelu', 'relu', 'softplus',
242                                     ''], default = 'tanh',
243                     help='activation function for the neural net')
243 parser.add_argument('-pr', '--problem', choices=['Bayesian', 'Generative'], default = 'Bayesian',
244                     help='What problem are we solving?')
245 parser.add_argument('-re', '--restore', type = int, default= 0,
246                     help='steps to restore')
247 args = parser.parse_args()
248
249 main(args)

```